24 May 1990

Technical Paper Sep 86 - May 90

The Expert Missile Maintenance Aid Program - Four Years Later and in the Field

F08635-86-C-0095 (C)
F08635-86-C-0393 (C)

Captain Barry Eugene Mullins, USAF

Raytheon Company
Missile Systems Division
Tewksbury, MA 01876-0904

Rockwell International
3370 Miraloma Avenue
Anaheim, CA 92803-3105

None

Air Force Armament Laboratory
AFATL/FXG
Eglin AFB, FL 32542-5434

AFATL-TP-90-13

Approved for public release; distribution is unlimited

The Expert Missile Maintenance Aid (EMMA) program is discussed in this paper. The paper addresses such issues as how tactical munition field-level and depot-level expert systems were developed and evaluated as well as the excellent results of the evaluations. It was found that the use of expert systems is very amenable and beneficial to diagnosing munition faults at both maintenance levels. Current work on the productization of the GBU-15 field-level system is briefly discussed as well as issues related to using DOD-STD-2167 to document the development of expert systems. When properly tailored, DOD-STD-2167 can be an effective means of documenting and monitoring the work performed in developing an expert system.

Artificial Intelligence, Expert Systems, Maintenance and Diagnostics, Integrated Diagnostics

14

Unclassified | Unclassified | Unclassified | SAR

THE EXPERT MISSILE MAINTENANCE AID PROGRAM - FOUR YEARS LATER AND IN THE FIELD

Barry E. Mullins, Captain, USAF
Air Force Armament Laboratory (AFATL/FXG)
Eglin Air Force Base, Florida 32542-5434

ABSTRACT

The Expert Missile Maintenance Aid (EMMA) program is discussed in this paper. The paper addresses such issues as how tactical munition field-level and depot-level expert systems were developed and evaluated as well as the excellent results of the evaluations. It was found that the use of expert systems is very amenable and beneficial to diagnosing munition faults at both maintenance levels. Current work on the productization of the GBU-15 field-level system is briefly discussed as well as issues related to using DOD-STD-2167 to document the development of expert systems. When properly tailored, DOD-STD-2167 can be an effective means of documenting and monitoring the work performed in developing an expert system.

INTRODUCTION

Current weapon systems are unquestionably profiting from technology advances and justifiably so. As a result, munitions are becoming more sophisticated, autonomous, and "smarter". Electronically sophisticated munitions are quickly infiltrating the Department of Defense (DOD) arsenal of weapons in order to combat the ever-increasing aggregate of sophisticated weapon systems of our adversaries. Simple bombs are becoming relics of yesteryears with technology driving our weapons towards autonomy.

Technology advances that improve the effectiveness of munitions are simultaneously complicating the maintenance of the same by increasing the munition's functionality typically at the expense of maintainability. Automatic test equipment (ATE) and associated test program sets (TPS) do not adequately diagnose munition faults. ATE (this paper refers to ATE as a collection of ATE and TPS) is plagued by high false alarm rates. Guidance and control sections returned to the depot are currently experiencing approximately 28 to 44 percent retest OKs. As many as one out of every four faults cannot be detected by ATE. Sequential testing (i.e., sequentially applying a series of predetermined tests to the faulty munition regardless of a priori knowledge) limits ATE diagnostic capabilities. Additionally, current munition ATE cannot diagnose beyond multiple linked components.

The current shortage of skilled munition maintenance technicians only exacerbates this dilemma. Demographic projections and the proposed cutbacks in DOD manpower forecast a deteriorating situation. Since an experienced technician is able to diagnose quicker and more reliably than a novice, the knowledge acquire by the experienced technician (expert) throughout the years should be captured so that this knowledge can be used by novice technicians during future diagnostic sessions.

Artificial intelligence (AI) technology is one approach to alleviating some of these problems while increasing the reliability and maintainability of existing and future weapon systems. Dr. Elaine Rich, University of Texas at Austin defines AI as follows: "Artificial intelligence is the study of how computers do things at which, at the moment, people are better" (Rich, 1983:1). Expert systems are a major subset of AI and have emerged recently with the greatest amount of success (Hayes-Roth et al., 1983:xi). Donald Waterman defines expert systems as "sophisticated computer programs that manipulate knowledge to solve problems efficiently and effectively in a narrow problem area" (Waterman, 1986:xvii). Tactical munition maintenance is one such domain in which a narrow problem area can be defined where a computer could potentially perform better than a human.

EMMA

EMMA is a research effort sponsored by the Air Force Armament Laboratory, Air Force Systems Command at Eglin Air Force Base (AFB), Florida. It is a first attempt to enhance maintenance of tactical munitions by employing AI. EMMA's ultimate goal is to assist a novice munition maintenance technician isolate and diagnose electronic, electromechanical, and mechanical faults of a munition to the board/chassis/component level more quickly and consistently than the best human expert using the best, currently-available ATE. To this end, EMMA augments existing ATE with an expert system that captures the knowledge of design and maintenance experts.

EMMA is a thirty month effort split into two phases which synchronize with the two levels of munition maintenance. Phase 1, September 1986 - July 1987, addressed the field-level maintenance of tactical munitions and ultimately resulted in two field-level expert systems. Phase 2, August 1987 - April 1989, focused on depot-level maintenance and likewise produced two depot-level expert systems. Phase 2 took twenty months (twice as long as phase 1) since depot-level diagnostic activities are more detailed than the field. The tactical munition maintenance technician is the intended user of EMMA. Since EMMA is constrained by schedule and money, the number of tests for each expert system developed under this effort is limited, yet sufficient, to demonstrate concept feasibility.

EMMA is a dual contract effort performed by Raytheon Company, Missile Systems Division in Tewksbury, Massachusetts, and Rockwell International Corporation, Autonetics Sensors and Aircraft Systems Division in Anaheim, California. Both contractors developed a field and depot EMMA resulting in a total of four systems (two for each contractor) and were allowed to select their candidate vehicle within specified limits. Raytheon selected the AIM-7F Sparrow missile as their munition; Rockwell chose the GBU-15 modular glide bomb.

EMMA draws on many different types of knowledge and information to perform the diagnosis of the faulty munition including maintenance rules or Technical Orders (TOs), maintenance technician practices (heuristics), Unit Under Test (UUT) design, existing test equipment capabilities, failure rates, and test costs. Figure 1 depicts how this knowledge is focused on

the problem of diagnosing the faulty munition. First, the symptoms are derived from the test equipment via a communication cable in most cases and technician observations via a user-friendly interface. The expert system then employs the knowledge stored in the knowledge bases and derives a repair strategy which is displayed to the technician using the EMMA computer screen.

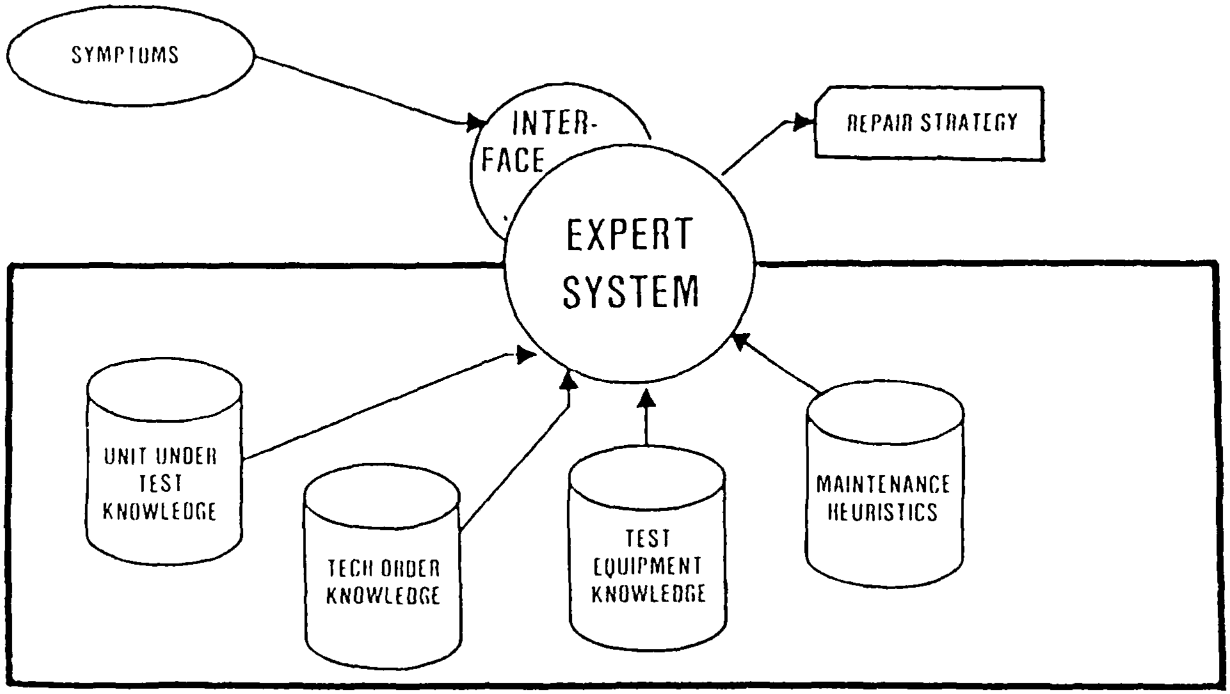

Figure 1. EMMA Expert System

PHASE 1 - THE AIM-7F FIELD-LEVEL EMMA

The Raytheon field-level EMMA is designed to enhance the field-level maintenance of the AIM-7F missile by augmenting the missile's test set -- AN/DSM-162. One of three AIM-7F subsystems is isolated in the field using the AN/DSM-162 -- target seeker, flight control, or tunnel cable. All references to the word "EMMA" in this section refer to the Raytheon AIM-7F EMMA. EMMA is hosted on a Symbolics 3670 LISP (LISt Processing) machine running the expert system shell ART (Automated Reasoning Tool). ART provides a production language that is primarily rule-based; consequently, EMMA uses the rule-based paradigm. The Symbolics computer is connected to the AN/DSM-162 via an RS-232 cable. Figure 2 illustrates the major components of the EMMA system and how they are interconnected.

The RS-232 cable allows EMMA to operate in one of three modes -- automatic, semi-automatic, and manual. The distinguishing characteristics of these modes is the level of automation EMMA is allowed during the diagnostic session. The automatic mode uses the RS-232 interface to allow EMMA to direct the diagnostic testing and resequencing of tests. EMMA

automatically accepts data from the test set via the RS-232 cable, performs the fault isolation, and directs the test set to perform additional tests, if required, until the fault is detected or all tests pass. If a fault is detected during automatic operation, the user may switch to semi-automatic mode for closer control over the testing and the ability to query after each test segment (i.e., use EMMA's explanation capability). Another advantage of the automatic modes (semi and full) is data integrity since EMMA passes data between the test set and the Symbolics computer and is not subject to operator errors. The manual mode is provided in the event an RS-232 connection is not possible and relies completely on a technician to shuttle information.

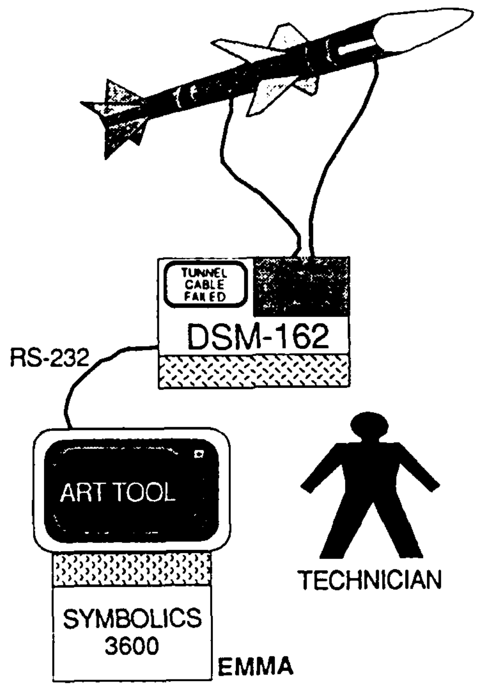

Figure 2. Field-level AIM-7F Configuration

As with most expert systems, EMMA is able to explain its reasoning process to the user (technician) by explaining its fault detection and resequencing logic in terms the technician understands (i.e., two explanation levels are available). The technician may request an explanation during any phase of the diagnostic process. This allows the technician to query EMMA during a consultation which heightens the technician's understanding of EMMA's activities while simultaneously providing the technician with a valuable training aid.

One of the most critical aspects of any software system is its user-friendliness. If the system is difficult to use and the user does not use it, it has failed. EMMA uses windows to relay information to the technician and accepts information via menus. Using a mouse, the

technician is able to enter data quickly and accurately without having to learn cryptic commands. The majority of the data entered into EMMA by the technician is done using the mouse; some keyboard input is required. Figure 3 shows the screen of a Symbolics computer running EMMA.

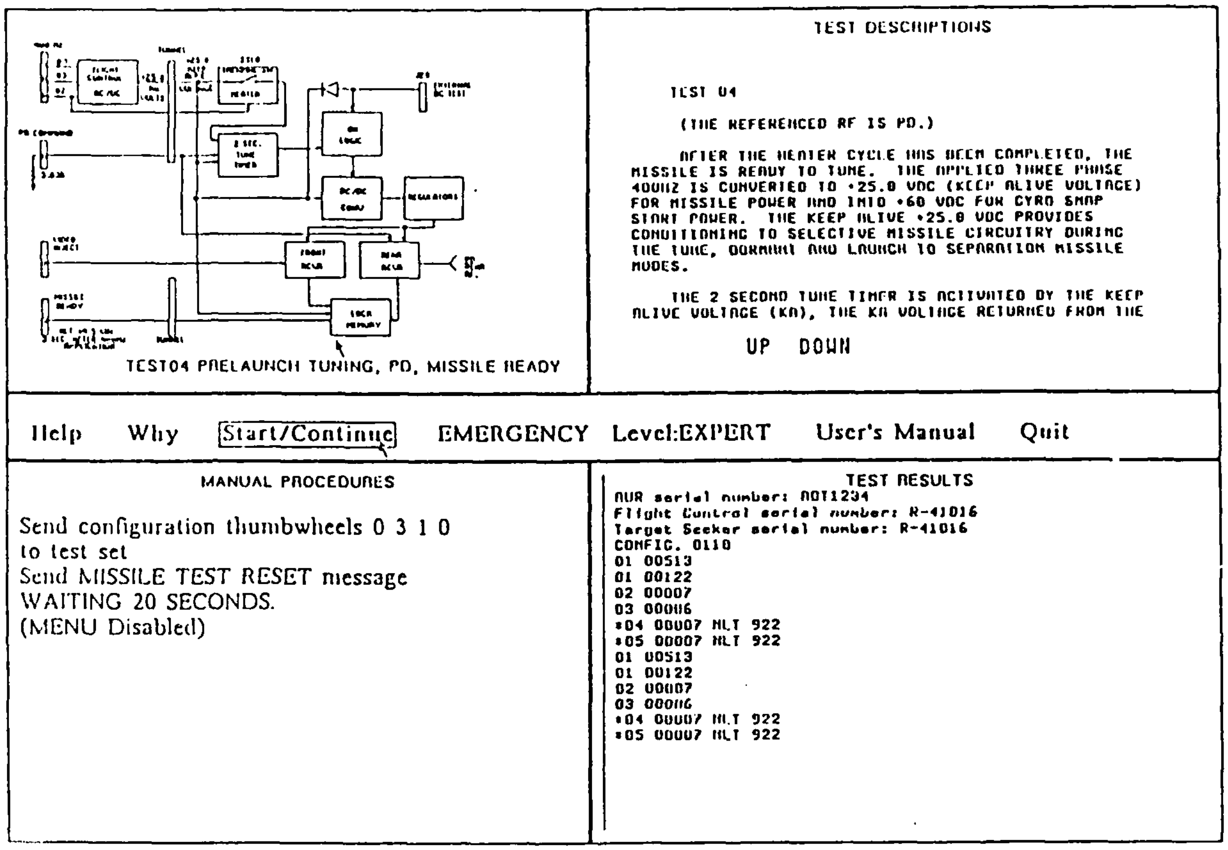

Figure 3. Field-level AIM-7F Screen

PHASE 1 - THE GBU-15 FIELD-LEVEL EMMA

The Rockwell field-level EMMA is designed to enhance the field-level maintenance of the GBU-15 by augmenting the field-level test set -- GJM-55. Depending on the proficiency of the technician, the GJM-55 can isolate to a subsystem such as the control module or to a circuit card. All references to the word "EMMA" in this section refer to Rockwell's GBU-15 EMMA unless stated otherwise. EMMA is hosted on an IBM PC/AT compatible computer running the M.1 expert system shell. Although, the M.1 language is primarily rule-based, EMMA uses an object oriented methodology. The rules of the knowledge base reference objects and object attributes. This EMMA does not support the capability for an automatic mode due to hardware limitations thereby leaving only the manual mode (i.e., no connecting cable). Figure 4 illustrates the major components of the EMMA system and how they are interconnected.

The GBU-15 EMMA also possesses explanation capabilities. Technicians may ask EMMA for an explanation or help at any time. EMMA responds with either an explanation of the reasoning process or information that guides the technician through the consultation. As with the AIM-7F EMMA, this EMMA has two levels of explanation to accommodate the needs of different technicians.

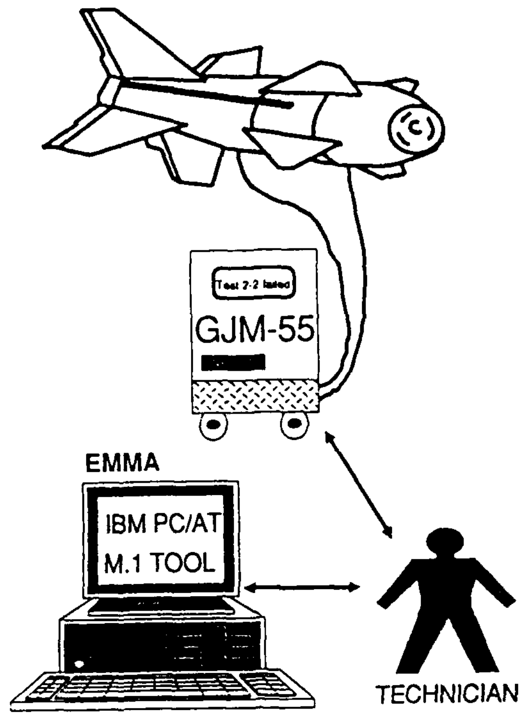

Figure 4. Field-level GBU-15 Configuration

EMMA exploits the use of pull-down menus and function keys to make it as user-friendly as possible. The majority of technician interaction with EMMA is performed using the keyboard. The technician typically responds to EMMA questions and requests with short answers thereby reducing the probability of erroneous data being entered.

EMMA EVALUATIONS

Meaningful evaluation of expert systems has been an often discussed but seldom achieved topic within recent years. More often than not quantitative metrics are simply not available or meaningful as an evaluation measure. Since an expert system encapsulates the knowledge of a given expert in a given field, the effective evaluation of the expert system may be difficult at best. Validation must be used to justify the representation levels of expert systems (O'Keefe et al., 1987).

Validation is typically considered a part of evaluation, and evaluation is concerned with determining the comprehensive value of an expert system (O'Keefe et al., 1987). Validation should not be confused with verification. "Validation refers to building the right system (that is, substantiating that a system performs with an acceptable level of accuracy), whereas verification refers to building the system 'right' (that is, substantiating that a system correctly implements its specifications)" (O'Keefe et al., 1987).

Verification of EMMA

A unique aspect of the verification of the EMMA program is it uses DOD-STD-2167, the Defense System Software Development standard, to develop the expert systems (field and depot). This is one of the first attempts to apply this standard to the development of an expert system. The EMMA program has shown that an expert system can be developed using DOD-STD-2167 software development requirements. With careful tailoring of some of the data item descriptions (DID), this standard can be effectively used to document the program and provide the program manager valuable insight into the contractor's software development, testing, and evaluation efforts. The tailored documents were slightly altered to accommodate the iterative nature of expert system development.

Since verification must determine whether an expert system correctly implements its specifications, testing must occur in order to validate this requirement. Again, DOD-STD-2167 proved to be adequate for verification testing once extended. Using the test documents in this standard, the correct implementations of the specifications for EMMA were verified. Two levels of testing occurred to accomplish this task. First, the knowledge engineer performed informal testing which verified the integrity of the individual computer software units before the units were integrated and tested as a system. Since expert system development is iterative in nature, informal testing essentially occurs throughout development when the knowledge engineer and expert verify the expert system's behavior. Second, an independent team performed formal testing by exercising EMMA using test plans and test descriptions generated using DOD-STD-2167. Both forms of testing identified problems which were later resolved.

Validation of EMMA

Validation of the field and depot systems is addressed in two areas: performance validation (i.e., how well EMMA performs), and the human factors aspect. Both areas are extremely important to the success of an expert system. The following paragraphs present the validation methodology and results of the two field-level EMMA systems. The depot EMMA validation results are presented later.

As with most expert systems, the ultimate measure of success is determined when it is used by the end users; the EMMA program is no exception. Both contractors took their respective field EMMAs to Air Force bases at which a field maintenance capability exists thereby allowing evaluation in an actual field test environment.

After considering several alternatives, both contractors decided to use a toggle switch box to simulate faults into a known good missile. This approach was necessary since it was feared that the maintenance locations might not have a sufficient number of faulty munitions/subsystems during the evaluation period. The faults were defined by the domain expert in conjunction with the maintenance expert (i.e., the experts selected various representative tests from the test set) such that the faults would adequately exercise the various characteristics of EMMA including the resequencing logic, the explanation capability, and the fault isolation logic.

The Field-level AIM-7F Evaluation. Raytheon took their field-level EMMA to the 325th Equipment Maintenance Squadron (EMS) at Tyndall AFB, Florida for an evaluation that began on 8 June 1987 and concluded on 12 June 1987.

Four EMS munition maintenance technicians were selected for the evaluation. Two technicians were novices with little AN/DSM-162 experience. The other two technicians were experts with substantial AN/DSM-162 experience. Two teams of two technicians were created consisting of one expert and one novice. One team (hereafter referred to as the EMMA team) received extensive EMMA training. The other team (hereafter referred to as the non-EMMA team) was not trained on the EMMA system and served as a baseline for the evaluation.

Twelve simulated faults were diagnosed by the EMMA team using EMMA and the non-EMMA team using just the AN/DSM-162 and the applicable TO. Performance of the two teams was based on the level of expertise of the operator, duration of test, and the ability to diagnose the fault accurately. The evaluation results were very promising. First, EMMA was able to consistently diagnose the fault quicker than the non-EMMA team using the AN/DSM-162 regardless of the experience level of the EMMA operator. A time savings of 20% was seen with the novice using EMMA versus the expert using the AN/DSM-162. Second, novice technicians using the EMMA system significantly outperformed (better fault diagnoses) novice and expert technicians using just the AN/DSM-162. Finally, EMMA's explanation capabilities significantly enhanced the abilities of the EMMA team to determine the reason behind each fault.

The argument could be made that EMMA should accurately diagnose all the induced faults since the expert system and the faults were derived from the same source -- the domain expert. In order to demonstrate the robustness of EMMA, an additional evaluation methodology was used. Two faulty missiles were saved by the EMS prior to the evaluation. These missiles had previously failed AN/DSM-162 testing. However, the fault data for these missiles were not released by the EMS personnel until after the EMMA evaluation. A third missile became available during the EMMA evaluation by failing a flight line test during prelaunch tuning. This missile was an excellent exercise for EMMA since its fault was unknown to everyone present. All three missiles (referred to as "mystery missiles" due to their unknown past) contained faults unknown to EMMA, the domain expert, and the maintenance technicians as well. EMMA was employed against these missiles with excellent results. The EMMA team correctly isolated

the faults in all three missiles. Only after EMMA diagnosed the faults were the previous test data on the two saved missiles released. EMMA's diagnosis was consistent with these data.

User acceptance was outstanding. In fact, the technicians accepted EMMA's diagnosis of the missile from the flight line and said, if allowed, they would have sent the missile to the depot with no further AN/DSM-162 testing. This exemplifies EMMA's acceptance by the EMS maintenance personnel who found EMMA to be very user-friendly. The mouse and the use of menus made the system easier to use than bulky and cumbersome TOs. Furthermore, the explanation capability proved to be an effective training mechanism.

The Field-level GBU-15 Evaluation. Rockwell evaluated their GBU-15 EMMA at the 4th EMS, Seymour Johnson AFB, North Carolina during the period of 22 June through 29 June 1987.

Once again, four maintenance technicians (two experts with several years experience and two novices with less than six months experience) were selected for the evaluation. All four technicians felt very comfortable using EMMA after a brief training session.

Twenty-two simulated faults were induced into the known good munition to evaluate EMMA's ability to handle the following five areas: resolution of ambiguities between major shop replaceable units (SRU), referencing lower configuration testing to facilitate further component resolution, distinguishing between a cable failure and a circuit card assembly (CCA) gain failure, resolution of ambiguities between CCA's, and recognizing operator errors or test set problems. Six of the twenty-two induced faults were in the all-up-round (AUR) configuration (i.e., the test was performed while the GBU-15 munition was completely intact). The remaining sixteen faults were in the control module stand alone configuration. EMMA was able to handle these five areas by analyzing additional test parameters as well as instituting and analyzing tests related to the failed test.

The diagnostic results of the induced faults revealed substantial time savings in fault isolation and increased diagnostic capabilities. While the munition was in the control module stand alone configuration, a time savings of 40% was realized over conventional testing with the GJM-55. When the munition was in AUR configuration, EMMA was able to provide up to 74% time savings. This is due to EMMA's capability to resolve failures while the munition is in the AUR configuration thereby saving the technician from having to performing testing in stand alone configuration.

The GJM-55, in some situations, will recommend more than one suspected failure -- an ambiguity group. EMMA was able to break up ambiguity groups. EMMA also considered the possibility of a cable harness or test set failure. Based on these capabilities EMMA was able to significantly improve fault isolation as seen by diagnosing the twenty-two simulated faults. EMMA added a wiring harness check to 50% of all tests. EMMA deleted a CCA from an ambiguity group 40% of the time thereby reducing the number of CCA to be considered during testing. EMMA added a CCA to an ambiguity group 30% of the time to insure all potential CCA's are

considered during the testing. This suggests that the test set did not always consider all potential CCA's. Finally, EMMA exchanged one suspect CCA in an ambiguity group for another CCA 10% of the time. The ability to manipulate the ambiguity group to facilitate fault isolation was demonstrated by EMMA and proved to be an effective fault isolation technique. These results directly support the time savings previously mentioned.

The GBU-15 EMMA also received kudos for its user-friendliness. The technicians used EMMA with comfort and found several items to be particularly laudable (e.g., EMMA's understandability). The explanation capability provided easy-to-understand responses. Another aspect they found beneficial was the addition of the internal wiring harness check as one of the reasons for a fault since this check is relatively "inexpensive" to perform and can prevent unnecessary and potentially costly future testing.

EMMA PHASE 2

Phase 2 is a natural extension of phase 1 since defective equipment found in the field are sent to the depot for repair. Due to the excellent results of phase 1, both contractors were allowed to proceed to phase 2 -- the depot-level maintenance of the same two munitions. The systems developed during phase 2 are more detailed extensions of the phase 1 systems with one exception; the phase 2 systems augment the depot-level test sets which are the AN/DSM-22 for the AIM-7F and the CATS (Computer Aided Test System) for the GBU-15.

The depot systems used the same methodology employed in phase 1. In fact, the same computer hardware and the same expert system shells are also used. However, the interfaces between the test sets and the EMMA computers are different. The Raytheon (AIM-7F) interface supports one-way communication from the AN/DSM-22 to the EMMA computer due to test set hardware limitations. The AN/DSM-22 isolates a circuit card (also known as a unit or tray) within the AIM-7F target seeker. Figure 5 illustrates the AIM-7F depot configuration. Rockwell (GBU-15) uses a two-way communication interface between the test set and the EMMA computer. The CATS tester facilitates the isolation of components (resistors, transistors, integrated circuits) on a GBU-15 circuit card. This depot configuration is shown in Figure 6. Both systems incorporate an explanation capability. The evaluation of the depot EMMAs also followed the same methodology used in phase 1. Each EMMA was evaluated at the actual depot location for the munition by the depot technicians.

THE DEPOT-LEVEL AIM-7F EVALUATION

The depot AIM-7F EMMA evaluation was conducted at the Missile Rework Facility (MRF), Naval Aviation Depot, NAS Alameda, California from 12 through 19 January 1989.

Two munition maintenance technicians (one novice and one expert) from the MRF were used for the evaluation. The novice technician had four years experience with the AN/DSM-22 but no experience with the AIM-7F target

seeker. The expert had eleven years of experience troubleshooting the AIM-7E and AIM-7F target seekers using the AN/DSM-22. Both received extensive EMMA training.

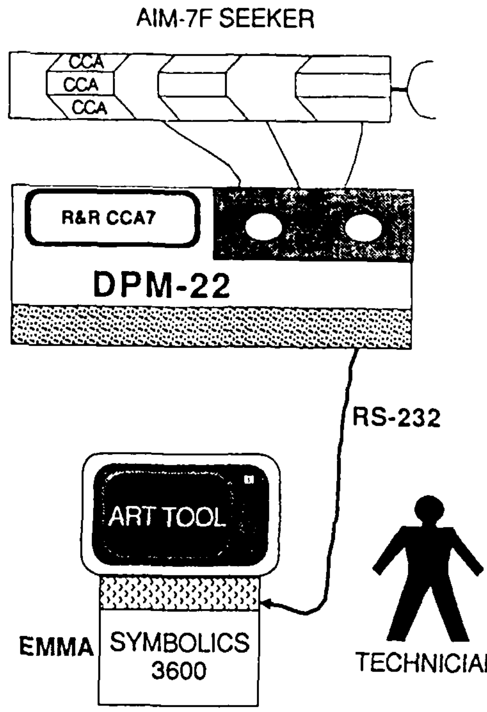

Figure 5. Depot-level AIM-7F Configuration

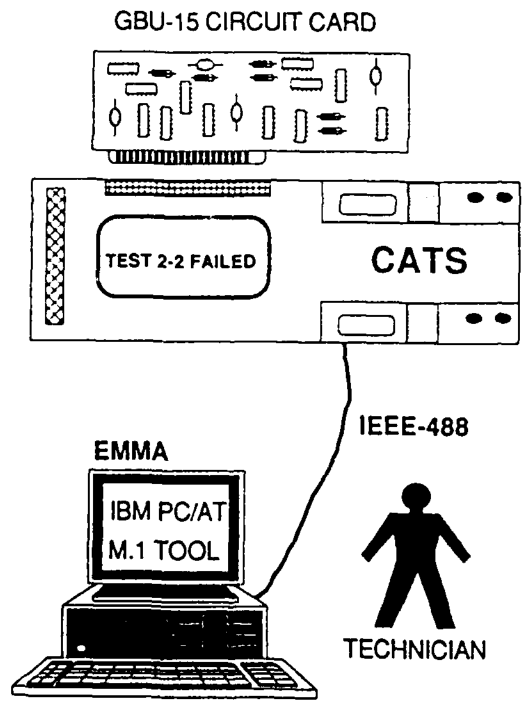

Figure 6. Depot-level GBU-15 Configuration

Five faults were inserted into a known good target seeker using jumper cables. The inserted faults were diagnosed by the novice using the EMMA system and then using the AN/DSM-22 and applicable procedures. The expert assisted by performing adjustments to the target seeker when requested by the test set for both evaluation runs (i.e., with and without EMMA). Although equally skilled, the novice was unfamiliar with the locations of the adjustments and would have spent extraneous time searching.

The evaluation results were excellent. The novice using EMMA was able to isolate 100% of the inserted faults 73% (average) quicker than the novice (same person) using just the AN/DSM-22 which found only 60% of the inserted faults. EMMA's explanation capabilities also enhanced the novice's ability to determine the reason behind each fault.

EMMA's abilities were further demonstrated by isolating four unknown (mystery) faults. EMMA successfully isolated three of the four faults. The one fault that EMMA was unable to isolate was a unique failure and not seen previously by the experts. However, using Raytheon's innovative "special problem" capability, the new fault was added to the knowledge base within 10 minutes.

THE DEPOT-LEVEL GBU-15 EVALUATION

Rockwell evaluated their GBU-15 EMMA at their Missile Systems Division in Duluth, Georgia since an organic depot capability currently does not exist. The evaluation occurred during the period of 13 through 17 February 1989. Figure 7 is an illustration of the color image seen by the GBU-15 technicians using EMMA.

Two munition maintenance technicians (one novice and one expert) were used for the evaluation. The novice technician had experience with GBU-15 diagnosis several years ago but no experience with the CATS tester nor GBU-15 circuit cards. The expert had five years experience troubleshooting GBU-15 circuit cards using CATS. Both technicians received EMMA training.

Sixteen faults were inserted into known good circuit cards using jumper cables. The inserted faults were then diagnosed by the novice using the EMMA system and then by the expert using CATS. The results of this evaluation were outstanding. The novice using EMMA was able to isolate 100% of the inserted faults 4 times faster (average) than the expert using CATS who also found all the faults. EMMA's real benefit was realized when the novice technician stated he would have been unable to isolate any of the faults without EMMA.

WHAT THE FUTURE HOLDS

EMMA has demonstrated the feasibility of applying AI to munition maintenance. In fact, due to the outstanding results of the field evaluations, Rockwell is converting the GBU-15 field-level system into a usable product under the sponsorship of the Armament Laboratory. Rockwell initially delivered the limited (number of GJM-55 testing configurations), development system to the 4th EMS, Seymour Johnson AFB in October 1989.

While the maintenance technicians are using/evaluating EMMA during their daily activities, Rockwell is continually updating EMMA to include more GJM-55 testing configurations, technician recommendations, graphics interfaces, and automatic generation of forms. The updated EMMA system is delivered to the 4th about every two months. The finished product will be available for deployment to TAC's GBU-15 field locations in September 1990. Funding for this completion effort was provided by the Productivity, Reliability, Availability, and Maintainability (PRAM) program office.

In future weapon systems, diagnostic expert systems should evolve with the weapon system instead of after the fact. This would allow the expert system to capture knowledge about the weapon as it is developed. Furthermore, the expert system should be incorporated/embedded directly into the ATE instead of augmenting the ATE with a separate computer system. This should make weapon systems of the future more supportable and maintainable by considering the maintenance aspect early in the weapon life cycle.

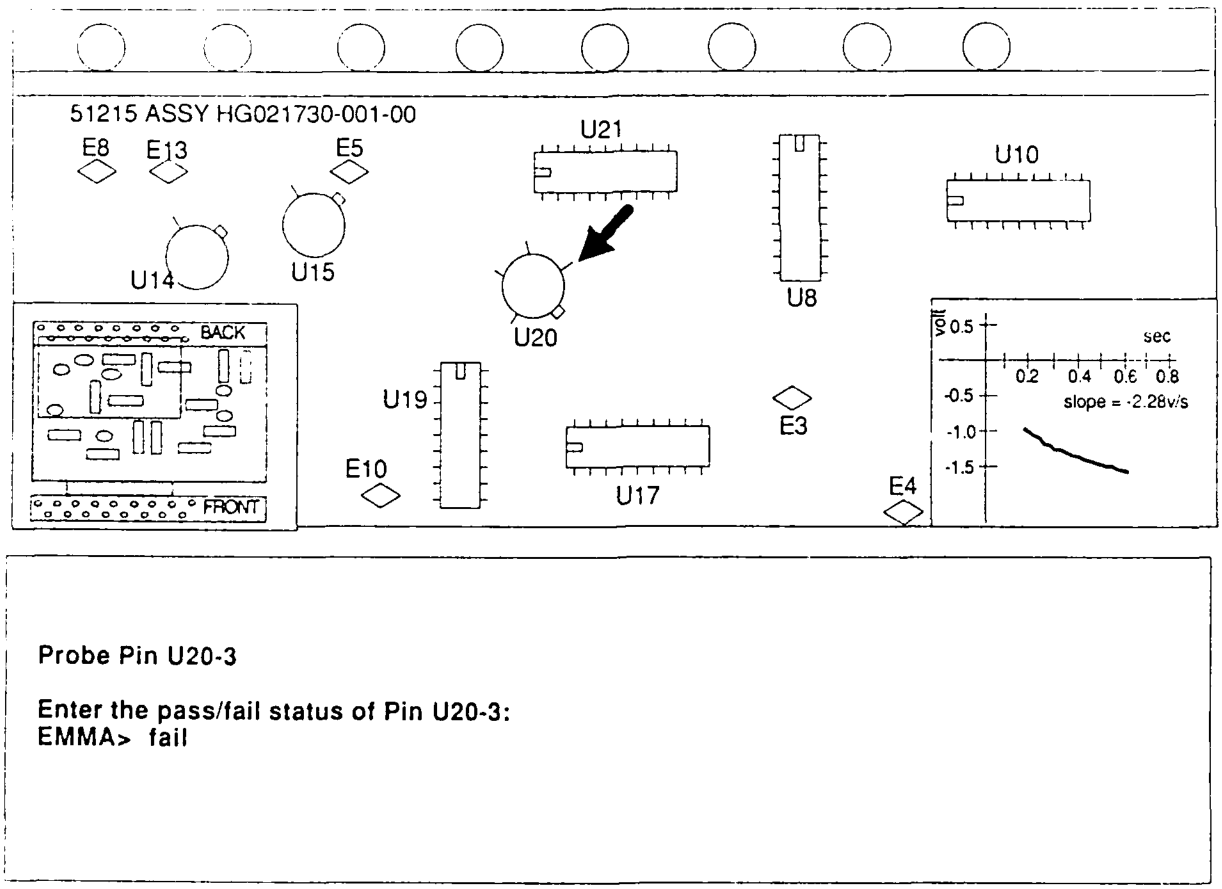

Figure 7. Depot-level GBU-15 Screen

SUMMARY

Tactical munition maintenance of today has problems. EMMA is an attempt to relieve some of these problems by applying expert system technology. EMMA's results reveal that this approach to munition maintenance has significant potential for the future.

Corporate knowledge retention is one of the premium benefits of EMMA. Since EMMA is updated easily and it never "forgets" knowledge, EMMA is an excellent tool for storing corporate knowledge as technicians come and go. Also, EMMA provides consistent, high quality diagnosis since it never has a "bad" day as contrasted with technicians. Rapid fault isolation and efficient manpower utilization are two more benefits of using EMMA. These benefits will result in substantial mission payoffs. Weapon system downtime will be decreased as well as personnel requirements and training time. However, the most significant payoff is the increase in the reliability of munition maintenance procedures and practices which will increase the overall readiness of the Air Force to provide hardware when and where required.

BIBLIOGRAPHY

Hayes-Roth, Fredrick et al. Building Expert Systems. Reading Massachusetts: Addison-Wesley Publishing Company, 1983.

O'Keefe, Robert M. et al. "Validating Expert System Performance," IEEE Expert, 2: 81-89 (Winter 1987).

Rich, Elaine. Artificial Intelligence. New York: McGraw-Hill Book Company, 1983.

Waterman, Donald A. A Guide to Expert Systems. Reading Massachusetts: Addison-Wesley Publishing Company, 1986.